

Sparse ℓ_1 Regularisation of Matrix Valued Models for Acoustic Source Characterisation

Laurent Hoeltgen · Michael Breuß ·
Gert Herold · Ennes Sarradj

Received: date / Accepted: date

Abstract We present a strategy for the recovery of a sparse solution of a common problem in acoustic engineering, which is the reconstruction of sound source levels and locations applying microphone array measurements. The considered task bears similarities to the basis pursuit formalism but also relies on additional model assumptions that are challenging from a mathematical point of view. Our approach reformulates the original task as a convex optimisation model. The sought solution shall be a matrix with a certain desired structure. We enforce this structure through additional constraints. By combining popular splitting algorithms and matrix differential theory in a novel framework we obtain a numerically efficient strategy. Besides a thorough theoretical consideration we also provide an experimental setup that certifies the usability of our strategy. Finally, we also address practical issues, such as the handling of inaccuracies in the measurement and corruption of the given data. We provide a post processing step that is capable of yielding an almost perfect solution in such circumstances.

Keywords Convex Optimisation · Sparse Recovery · Split Bregman · Matrix Differentiation · Acoustic Source Characterisation · Microphone Array

PACS PACS 43. · PACS 43.20.-f

Mathematics Subject Classification (2010) MSC 65K10 · MSC 65Z05 · MSC47A99

This work has partially been funded by the German Research Foundation (DFG) within the grant number SA 1502/5-1. This funding is thankfully acknowledged.

L. Hoeltgen · M. Breuß
Chair for Applied Mathematics, Brandenburg University of Technology Cottbus-Senftenberg
Platz der Deutschen Einheit 1, 03046 Cottbus, Germany
E-mail: {hoeltgen, breuss}@b-tu.de

G. Herold · E. Sarradj
Chair for Technical Acoustics, Brandenburg University of Technology Cottbus-Senftenberg
Siemens-Halske-Ring 14, 03046 Cottbus, Germany
E-mail: {herold, sarradj}@b-tu.de

1 Introduction

In 2005, Osher *et al.* [31] proposed an algorithm for the iterative regularisation of inverse problems that was based on findings of Bregman [3]. They used this algorithm, nowadays called Bregman iteration, for image restoration purposes such as denoising and deblurring. Especially in combination with the Rudin-Osher-Fatemi (ROF) model for denoising [37] they were able to produce excellent results. Their findings caused a subsequent surge of interest in the Bregman iteration. Among the numerous application fields, it has for example been used to solve the basis pursuit problem [8, 32, 50] and was later applied to medical imaging problems [16]. Further applications include deconvolution and sparse reconstructions [51], wavelet based denoising [45] and nonlinear inverse scale space methods [5, 6]. An important adaptation of the Bregman iteration is the Split Bregman (SB) method [14] and the linearised Bregman approach [8]. The SB algorithm can be used to solve ℓ_1 -regularised inverse problems in an efficient way. Its benefits stem from the fact that differentiability is not a necessary requirement on the underlying model.

Our contribution is concerned with an application to acoustic source characterisation using a microphone array. A microphone array comprises n microphones at known locations. These microphones register the sound that is emitted by a number of sources with unknown locations (see Fig. 1). The characterisation of these acoustic sources requires the estimation of their location and strength.

We briefly describe the physical background behind the models that we discuss in this work. The propagation of sound from a source position x to a receiver at position y can be modelled by a Green's function. In the following we assume that the source is always a monopole. In that case the sound pressure amplitude at the receiver position for a given discrete frequency ω is defined by

$$p(r, \omega) = q_0 \frac{1}{r} \exp\left(-\iota \omega \frac{r}{c_0}\right) \quad (1)$$

with ι being the complex unit, q_0 being the source strength, and r denoting the distance between the source and the receiver. Finally, the constant c_0 denotes the speed of sound. The signals from any given source are evaluated at a reference point y_0 at distance r_0 from the source:

$$p_0(r_0, \omega) = q_0 \frac{1}{r_0} \exp\left(-\iota \omega \frac{r_0}{c_0}\right) \quad (2)$$

Introducing the reference point into our model formulation helps us to eliminate the source strength and leads to the following description

$$p(r, \omega) = a p_0(r_0, \omega) \quad (3)$$

with

$$a := \frac{r_0}{r} \exp\left(\iota \omega \frac{r_0 - r}{c_0}\right) \quad (4)$$

Equation (3) yields the sound pressure amplitude at a receiver position depending on the sound pressure induced at a reference location by the source.




Fig. 1 A generic set-up for acoustic source characterisation using a microphone array. Each source s causes a sound pressure amplitude x_s at the reference point. The sound pressure c_j arriving at microphone j is given by $c_j = a_{sj}x_s$. If multiple sources are present, then it follows from a superposition principle that we have $c_j = \sum_i a_{ij}x_i$.

The estimation approach that we follow here additionally assumes that the actual locations of the sources are restricted to m possible source locations. Since a superposition principle holds in our model, we can account for multiple sources by adding all contributions. Thus, the sound pressure at a microphone j is given by $\sum_i a_{ij}x_i$ where the sum is taken over all possible source locations i . The coefficient a_{ij} is defined in accordance with (4):

$$a_{ij} := \frac{r_{0,j}}{r_{ij}} \exp\left(i\omega \frac{r_{0,j} - r_{ij}}{c_0}\right) \quad i = 1, \dots, m, \text{ and } j = 1, \dots, n \quad (5)$$

with r_{ij} denoting the distance between sender i and receiver j .

In a typical setting the number m of possible source locations is much larger than the number n of microphones, but the number of actual sources is less than n . Therefore, most of the x_i in the sum are zero. The described setup is also visualised in Fig.1.

Gathering all possible coefficients a_{ij} in a matrix A , and using a vector $c \in \mathbb{C}^n$ to hold all microphone sound pressures yields

$$c = Ax, \quad (6)$$

where $x \in \mathbb{C}^m$ is the sparse vector of source strengths. Using the Hermitian cross-spectral matrix $C := E[cc^\top]$ of microphone sound pressures, we can reformulate (6) as

$$C = A E \underbrace{\begin{bmatrix} xx^\top \end{bmatrix}}_{=X} A^\top \quad (7)$$

where the operator E denotes the expected value and where $X \in \mathbb{C}$ is the cross spectral matrix of source levels. This matrix is sparse, Hermitian, and in case of uncorrelated source signals also diagonal. By estimating X , the task of characterising the sources is solved. Let us emphasise that this is an ill-posed inverse problem. We refer the reader to [17–19] for further details.

Our Contribution. In this paper we discuss different models to characterise sound sources. Our starting point is (7). Instead of solving this linear system of equations directly we suggest to minimise the corresponding least squares formulation and to augment our model with a sparsity inducing regulariser. Two different approaches can be taken into account to respect the assumption that the solution X from (7) should be a sparse diagonal matrix. We present a formulation that adds harsh penalisation weights to off-diagonal entries, and a second formulation that seeks optimal solutions over the space of diagonal matrices. Both approaches use ℓ_1 norms of the solution X to enhance sparsity.

Further, we propose numerical strategies for both models. Our basic algorithm is an adapted version of the popular SB method, which has been used successfully for sparsity constrained problems in the past. It is well suited to handle the non-differentiability of our cost function. For our second model we additionally introduce advanced matrix valued differentiation concepts to preserve the diagonal structure of the solution. We see the main novelties of our work in these two ideas, namely the application of the SB scheme on the task at hand as well as the usage of matrix differential calculus to enforce structural constraints in the solution.

Finally, we present a post processing strategy to handle setups with corrupted data. The benefits of our approach are a robust behaviour with very sparse solutions. Our solvers are tailored such that the number of non-zero entries in their solutions coincides exactly with the number of sought sources.

Structure of the Paper. In Section 2 the considered problem is discussed in more detail. We elaborate on the occurring difficulties and propose a certain number of models to overcome the mentioned difficulties in Section 2.1. Section 3 presents our considered numerical approaches. The necessary mathematical preliminaries are briefly outlined in Section 3.1 whereas Sections 3.2-3.4 provide a detailed presentation of our numerical schemes. Potential post processing methods to improve the results further are established in Section 3.5. Afterwards we provide a thorough evaluation on real world data sets in Section 4. Finally, we give a short conclusion with an outlook on challenges and future work in Section 5.

2 Problem Formulation

The mathematical formulation of the considered problem is as follows. Given matrices $A \in \mathbb{C}^{n,m}$, and $C \in \mathbb{C}^{n,n}$, where C is also known to be Hermitian, we seek a diagonal matrix $X \in \mathbb{C}^{m,m}$ with a sparse set of non-zero entries along its main diagonal such that the equation

$$AXA^\top = C \quad (8)$$

holds. We remark that the previous equation can easily be rewritten as

$$(A \otimes A) \text{vec}(X) = \text{vec}(C) \quad (9)$$

where \otimes represents the Kronecker matrix product and where the vec operator stacks the entries from its argument column wise on top of each other. We refer to [25, 28, 33]

for a thorough presentation and properties of these two operators. It follows that one could actually reduce the initial problem in (8) to solving the linear system given in (9). However, the system matrix is dense and prohibitively large and thus unsuited for most computational routines. If each entry in $A \otimes A$ is stored in double precision, then we require $16n^2m^2$ byte of memory. For the experimental setups that we consider in this paper this equals to about 185 GB of data. In addition, we cannot assert that all our model assumptions will always be fulfilled if we simply solve the linear system. Indeed, in practice the matrices A and C will stem from experimental setups and be corrupted with noise. This means, even if we were to find a solution to (8), we cannot expect it to be a real valued diagonal matrix with a sparse set of non-zero entries along its main diagonal. Indeed, we have the following result, which can also be found verbatim in [25] (Chapter 2, Theorem 3)

Proposition 1 *A necessary and sufficient condition for the matrix equation $AXB = C$ to have a solution is that $AA^\dagger CB^\dagger B = C$, in which case the general solution is*

$$X = A^\dagger CB^\dagger + Q - A^\dagger AQB B^\dagger \quad (10)$$

where Q is an arbitrary matrix of appropriate order.

Here, A^\dagger denotes the Moore-Penrose inverse of the matrix A . Applying the previous proposition onto (8) with $B = A^\top$ yields conditions that assert the existence of a solution. Also, it becomes obvious that the solution of the system may not necessarily have the sought structure.

In order to overcome the aforementioned restrictions we suggest to consider the following variational models

$$\arg \min_{X \in \mathbb{C}^{m,m}} \left\{ \frac{1}{2} \|AXA^\top - C\|_F^2 + \lambda \|X\|_1 \right\} \quad (11)$$

$$\arg \min_{X \in \mathbb{C}^{m,m}} \left\{ \frac{1}{2} \|AXA^\top - C\|_F^2 + \|W \circ X\|_1 \right\} \quad (12)$$

$$\left\{ \begin{array}{l} \arg \min_{X \in \mathbb{C}^{m,m}} \left\{ \frac{1}{2} \|AXA^\top - C\|_F^2 + \lambda \|X\|_1 \right\} \\ \text{under the constraint that } X \text{ is a diagonal matrix} \end{array} \right\} \quad (13)$$

where \circ represents the Hadamard matrix product [27]. It simply multiplies all matrix entries componentwise with each other. Also, $\|\cdot\|_F$ represents the Frobenius norm for matrices [20, 21]. In order to treat complex valued arguments, we additionally use the following conventions:

$$\|X\|_1 := \sum_{i,j} (|\operatorname{Re}(X_{ij})| + |\operatorname{Im}(X_{ij})|) \quad (14)$$

$$\|X\|_F^2 := \sum_{i,j} (\operatorname{Re}(X_{ij})^2 + \operatorname{Im}(X_{ij})^2) \quad (15)$$

Here, $\operatorname{Re}(z)$ and $\operatorname{Im}(z)$ represent the real and imaginary part of the complex number z . Our choice is motivated by the fact that it will allow us a fast parallel optimisation in the construction of our numerical schemes. An in-depth analysis on the models is

given in the forthcoming Section 2.1 and the exact details on the benefits of employing (14) and (15) will be discussed in Section 3.2.

Let us also remark that the considered task in obtaining a representation AXA^\top from a given matrix C bears strong similarities with matrix factorisation problems, as described in [24, 46, 47], where models, similar to ours, are also commonly used. Further, our models bear a certain resemblance to Basis Pursuit (BP) approaches, too. These usually seek solutions of

$$\arg \min_{x \in \mathbb{R}^n} \left\{ \frac{1}{2} \|Bx - p\|_2^2 + \mu \|x\|_1 \right\} \quad (16)$$

for some given data $B \in \mathbb{R}^{m,n}$, $p \in \mathbb{R}^m$ and a positive regularisation weight μ . Such models are for example discussed in [17, 18]. Possible algorithms to obtain a solution to (16) include Lasso [42], Forward Backward Splitting (FBS) [39, 40], SB [14] and Primal Dual Hybrid Gradient (PDHG) [10, 34]. The PDHG algorithm requires an estimate of the eigenvalue spectrum of the matrix B in order to set the initialisation correctly. Depending on the size and the structure of the matrix, such an estimate may be difficult to get. The FBS can be very fast if additional techniques to optimise the parameters of the framework are being used. Split Bregman is very robust and does not need any special parameter setups. This observation as well as its other convenient theoretical properties make SB our algorithm of choice to base our numerical strategies on.

2.1 Analysis of the Suggested Models

Let us analyse the suggested models and emphasise their properties in more detail. The models proposed in (11)-(13) follow a hierarchy in increasing complexity and accuracy on our model assumptions. All three models are composed of two terms. The first term being a least squares formulation of our assumption that (8) should be fulfilled. The second term steers the sparsity in the solution. The ℓ_1 penaliser is a common choice to enforce only few entries to be non-zero and easier to handle from an optimisation point of view than directly penalising the number of non-zero entries. The latter option would lead to a non-continuous combinatorial problem. The cost function of each proposed model is convex, but not necessarily strictly convex. Furthermore, only (11) and (13) possess cost functionals that are always coercive.

The formulation in (11) certainly resembles canonical ℓ_1 regularised least-squares models most. The real valued non-negative scalar λ acts as a regularisation weight and influences the sparsity of the solution. Larger parameter choices yield sparser solutions without allowing any precise specification of the sparsity pattern. Its minimisation will be the easiest to carry out but we expect the quality of the results to be inferior to the other models, too. A diagonal matrix with only non-zero entries on its main diagonal can already be regarded as sparse when compared to an arbitrary matrix. Yet, we wish to have few non-zero entries along the main diagonal only. Thus the model does not fully coincide with our assumptions.

Equation (12) allows a very fine grained tuning of the sparsity structure. The weighting matrix W with non-negative entries fulfils a similar purpose as the parameter

λ in (11) but allows a more differentiated weighting of the individual entries. Setting all off-diagonal entries to large values will favour solutions with diagonal structure. However, a sparse structure along the main diagonal will still be difficult to achieve. Suitable choices of the weighting matrix W could also enable certain non diagonal entries to be positive. Such set-ups would allow the handling of correlated acoustic sources and could be the subject of future investigations. The challenge lies in finding suitable weight distributions to accurately model the correlation between individual sources. If all entries in W are identical, then (12) coincides with (11).

Lastly, (13) comes closest to our model assumptions that X should be a sparse diagonal matrix. The structural constraint is built explicitly into the model. Such constrained optimisation tasks are usually more challenging than their unconstrained counterparts, but in this case the cost considered in (13) comes closest to our model assumptions and is likely to yield the best results. This model can be seen as the limiting case when all off-diagonal elements of the weighting matrix W in (12) tend to $+\infty$.

Since (11) represents a well studied model from the literature and since (12) and (13) fit better to our model assumptions we will mostly focus on the latter two formulations in this work.

3 Our Novel Solution Strategy

All suggested models are structured like

$$\arg \min_x \{f(x) + g(Ax - b)\} \quad (17)$$

for some functions f , g , a matrix A , and a vector b . We note that this setup also includes constrained optimisation tasks. It suffices to set $g(x) = \iota_S(x)$, where ι_S is the indicator function of the set S to force the solution to be inside the set S . Tasks like (17) are well suited to be solved by splitting schemes. Among the vast choice of existing schemes we name the already mentioned FBS, PDHG, SB as the most popular ones. The FBS algorithm probably belongs to the best studied approaches and several extensions exist to further improve its efficiency. The PDHG approach excels in terms of speed, especially if it is used in conjunction with preconditioning techniques such as those presented in [34]. The SB method is a very viable strategy with a thorough convergence theory and sufficient flexibility to be applied to a large number of distinct formulations.

In this work, we will adapt the SB approach to our setting. Our choice is motivated by the fact that the SB method can be formulated such that it applies to all our models and that it requires only a single parameter, which in addition has a very intuitive interpretation.

3.1 Mathematical Preliminaries

In this section we briefly regroup a certain number of findings from the literature that we will rely on in the forthcoming sections. These results stem mostly from matrix (differential) calculus. We refer to [20, 21, 25, 28, 33, 35] for an in-depth discussion.

3.1.1 Matrix Differential Theory

The following results show that the space of complex valued matrices bears a Hilbert space structure when equipped with the Frobenius norm. This observation allows us to simplify certain expressions in the forthcoming sections and provide a convenient framework to operate in.

Definition 1 (Matrix Scalar Product [33]) Let $A = (a_{ij})$, $B = (b_{ij})$ be two $m \times n$ matrices over \mathbb{C} . We call *matrix scalar product* the expression $\langle A, B \rangle$ defined by

$$\langle A, B \rangle := \text{tr}(A^\top B) = \sum_{i=1}^m \sum_{j=1}^n \overline{a_{ij}} b_{ij} . \quad (18)$$

Here, the operator $\text{tr}(\cdot)$ denotes the trace of a matrix. The matrix scalar product is often called *Frobenius scalar product* because it induces the Frobenius norm.

Lemma 1 Let $A = (a_{ij})$ be a $m \times n$ matrix over \mathbb{C} . The Frobenius norm is related to the matrix scalar product in the following way:

$$\|A\|_F^2 := \sum_{i=1}^m \sum_{j=1}^n |a_{ij}|^2 = \sum_{i=1}^m \sum_{j=1}^n \overline{a_{ij}} a_{ij} = \langle A, A \rangle \quad (19)$$

Proposition 2 The matrix scalar product is a scalar product in the proper mathematical sense, i.e. all properties that are known to hold for general scalar products also apply to the matrix scalar product. In particular, we have the following rules for complex valued matrices A , B , and C :

$$\|A \pm B\|_F^2 = \|A\|_F^2 \pm 2\langle A, B \rangle + \|B\|_F^2 \quad (20)$$

$$\langle AB, C \rangle = \langle B, A^\top C \rangle \quad \text{and} \quad \langle A, BC \rangle = \langle B^\top A, C \rangle \quad (21)$$

Proof The proofs are straightforward. Equation (20) follows from the bilinearity of the scalar product. In order to show (21) it is helpful to use the following result from [33] (Equation (16)):

$$\text{tr}(ABC) = \text{tr}(BCA) = \text{tr}(CAB) \quad (22)$$

□

Splitting schemes such as the SB method often assume that the cost function to be minimised can be decomposed into simpler (convex) terms which will be optimised in certain ways alternatingly. For performance reasons it is beneficial if these optima have closed form solutions. Such representations can usually be derived from first order optimality conditions. The next propositions state a certain number of findings related to matrix valued derivatives that will prove to be helpful. A thorough discussion on the notion of matrix differentials and their derivatives can be found in [25]. As we mainly make use of the results documented in the aforementioned references we only state the necessary formulas here.

Corollary 1 We have $\langle AXB^\top, Y \rangle = \langle X, A^\top YB \rangle$.

Proof Follows from (22).

Computing derivatives with respect to matrices requires to take the structure of the matrix into account. In a symmetric matrix off-diagonal entries appear always twice for example. We denote the derivative with respect to a matrix X with $\frac{d}{dX}$. If X is assumed to have a specific structure we will mention it explicitly.

Proposition 3 *We have the following identities for an arbitrary unstructured matrix X :*

$$\frac{d}{dX} \operatorname{tr} \left((AXB + C)^\top (AXB + C) \right) = \frac{d}{dX} \|AXB + C\|_F^2 \quad (23)$$

$$= 2A^\top (AXB + C)B^\top \quad (24)$$

$$\frac{d}{dX} \|W \circ X - C\|_F^2 = 2W \circ (W \circ X - C) \quad (25)$$

From (23) and (24) we may conclude that necessary optimality conditions for a minimiser of $\frac{1}{2} \|AXA + C\|_F^2$ are given by $A^\top (AXA + C)A^\top = 0$.

Proof Eq. (24) is stated explicitly in [33] as Eq. (108), whereas (25) follows from the chain rule (Theorem 12, Chapter 5 in [25]) and the differential of the Hadamard product (Eq. (17), Chapter 8 in [25]).

If we require that the matrix is structured, then other identities hold [28]. When working with symmetric (resp. diagonal) matrices it seems intuitive to expect derivatives to yield symmetric (resp. diagonal) matrices as well. Especially, if we restrict our attention to diagonal matrices, then we obtain the following equality.

Proposition 4 *We have the following equality for a diagonal matrix X :*

$$\frac{d}{dX} \frac{1}{2} \|AXB - C\|_F^2 = \left(A^\top (AXB - C)B^\top \right) \circ I \quad (26)$$

where I is the identity matrix.

Proof Follows from the formulas established in Proposition 3 under consideration of the restrictions stated in [33] (Section 2.8).

By comparing (23) and (26) we clearly see the importance of taking the structure of the underlying matrices into account. The derivative in (26) only coincides with the derivative in (23) along the main diagonal. A similar finding can be deduced for symmetric matrices, too.

3.1.2 Matrix Valued Proximal Operators

As already mentioned, splitting algorithms decompose the original task into simpler subproblems. Some of these subtasks coincide with proximal operations which go back to Moreau [29]. We also refer to [11] for a detailed analysis of their properties. The proximal operator of the ℓ_1 norm will play a significant role in our framework. It is a well known operator for signal processing tasks, where it is often referred to as soft shrinkage.

Definition 2 (Soft Shrinkage) The soft shrinkage operator $\text{shrink}_\lambda(b)$ with parameter $\lambda > 0$ solves the optimisation problem

$$\arg \min_{x \in \mathbb{R}^n} \left\{ \lambda \|x\|_1 + \frac{1}{2} \|x - b\|_2^2 \right\} \quad (27)$$

For real valued data the soft shrinkage operator has a closed form representation. We provide this result with a detailed proof since we could not find a source with a satisfying derivation.

Lemma 2 Let $b \in \mathbb{R}^n$. The solutions of (27) are given by

$$\begin{aligned} (\text{shrink}_\lambda(b))_i &:= \text{sgn}(b_i) \max(|b_i| - \lambda, 0) \\ &= \begin{cases} b_i + \lambda, & b_i < -\lambda \\ 0, & b_i \in [-\lambda, \lambda] \\ b_i - \lambda, & b_i > \lambda \end{cases} \end{aligned} \quad (28)$$

Proof First, we note that (27) decouples into n independent optimisation problems of the form

$$\arg \min_{x_i \in \mathbb{R}} \left\{ \lambda |x_i| + \frac{1}{2} (x_i - b_i)^2 \right\} \quad i = 1, \dots, n \quad (29)$$

Thus, it suffices to solve the simpler one dimensional optimisation task and to rearrange the result componentwise. We note that the energy in (29) is a strictly convex function and thus every local minimum must necessarily also be a global minimum. Minimisers \bar{x}_i must fulfil the first order optimality condition $0 \in \lambda \partial(|\cdot|)(\bar{x}_i) + (\bar{x}_i - b_i)$, where $\partial(|\cdot|)(\bar{x}_i)$ is the subdifferential of the absolute value function evaluated at \bar{x}_i . Let $q \in \partial(|\cdot|)(\bar{x}_i)$ be any subgradient, then we obtain

$$0 = \lambda q + \bar{x}_i - b \Leftrightarrow q = \lambda^{-1}b - \lambda^{-1}\bar{x}_i \quad (30)$$

There are three possible choices for q . If $\bar{x}_i > 0$ then $q = 1$, if $\bar{x}_i < 0$ then $q = -1$ and if $\bar{x}_i = 0$, then $q \in [-1, 1]$. Now, if $\lambda^{-1}b > 1$ then \bar{x}_i must be positive because q cannot be larger than 1. But for positive \bar{x}_i , $\partial(|\cdot|)(\bar{x}_i) = \{q\} = \{1\}$ and \bar{x}_i must be $b - \lambda$. On the other side, if $\lambda^{-1}b < -1$, then \bar{x}_i must be negative, since q cannot be smaller than -1 . For negative \bar{x}_i , $\partial(|\cdot|)(\bar{x}_i) = \{q\} = \{-1\}$ and \bar{x}_i is given by $b + \lambda$. Now assume $\lambda^{-1}b \in [0, 1]$, then $-\lambda^{-1}\bar{x}_i = q - \lambda^{-1}b$ must hold. If we suppose $q > \lambda^{-1}b$, then \bar{x}_i would have to be strictly negative. However, for strictly negative \bar{x}_i , $q = -1$ and we have a contradiction. In the same way we cannot have $q < \lambda^{-1}b$ since then we must have $\bar{x}_i > 0$ and $q = 1$. It follows that $q = \lambda^{-1}b$ is the only remaining choice. This implies $\bar{x}_i = 0$. The case $\lambda^{-1}b \in [-1, 0]$ leads by identical reasoning to the same result $\bar{x}_i = 0$ and $q = \lambda^{-1}b$. Regrouping all solutions completes the proof. \square

Lemma 2 yields a closed form description of the proximal mapping for real valued arguments only. However, our task at hand has complex valued data. Ignoring the imaginary parts in the computations would lead to significant errors. The following theorem shows how to extend the previous finding. We note, that it relies on our choice for the complex valued norms from (14) and (15). For other choices it might be difficult or even not possible at all to get such a closed form representation.

Theorem 1 *Let $B \in \mathbb{C}^{n,m}$. Then, we have*

$$\arg \min_{X \in \mathbb{C}^{n,m}} \left\{ \lambda \|X\|_1 + \frac{1}{2} \|X - B\|_F^2 \right\} = (\text{shrink}_\lambda(\text{Re}(B)) + \iota \cdot \text{shrink}_\lambda(\text{Im}(B)))_{ij} \quad (31)$$

with ι denoting the complex unit.

Proof Due to our convention in (14) and (15), the minimisation in (31) not only decouples into nm smaller optimisation problems, each in a single scalar valued complex unknown, but we can further optimise real and imaginary part independently:

$$\begin{aligned} \arg \min_{X \in \mathbb{C}^{n,m}} \left\{ \lambda \|X\|_1 + \frac{1}{2} \|X - B\|_F^2 \right\} = \\ \arg \min_{X \in \mathbb{C}^{n,m}} \left\{ \lambda \sum_{i,j} |\text{Re}(X_{ij})| + |\text{Im}(X_{ij})| + \frac{1}{2} (\text{Re}(X_{ij} - B_{ij})^2 + \text{Im}(X_{ij} - B_{ij})^2) \right\} \end{aligned} \quad (32)$$

Thus, we obtain $2mn$ optimisation problems that can be solved independently by making use of Lemma 2. \square

We note that the previous theorem naturally extends the case of real valued matrices as well as the setting with vector valued data. It also represents an important finding for our numerical scheme. The componentwise handling of the data allows an efficient scaling of our numerical solvers to large parallel processing facilities such as General Purpose Computation on Graphics Processing Units (GPGPUs).

3.2 The Bregman Framework

We present a short review of the (split) Bregman algorithm, which will present the basis of our forthcoming developments. There exists a large family of algorithms that could be considered to form the Bregman framework. The most prominent representatives would be the *standard Bregman iteration*, as developed by Osher and colleagues [31], the *linearised Bregman algorithm* by Cai et al. [8], and the *split Bregman method* by Goldstein and Osher [14]. Bregman himself [3] originally wanted to describe non-orthogonal projections onto convex sets and derived, as by-product, an iterative scheme to minimise certain smooth and convex functions under linear constraints. We refer to [1] for a deeper analysis of these projections. Let us also remark that the Bregman algorithms are related to many other popular strategies from the literature and a certain number of equivalences with proximal methods have been discovered in the past. An extensive discussion on these relationships can for example be found in [40]. Convergence considerations for various setups can, in addition to the already cited works, also be found in [4, 5, 7, 9, 14]. Finally, various strategies to improve the computational performance of the split Bregman scheme has been discussed in [12, 13, 15, 23] and the references therein.

3.2.1 The Split Bregman Method

The SB method has originally been developed to solve convex ℓ_1 penalised optimisation tasks. In its simplest form it aims to find a solution of equations such as (17) with convex but not necessarily smooth functions f and g . The underlying idea is to introduce a slack variable z and to reformulate the original task in the equivalent form

$$\begin{aligned} & \arg \min_{x, z} \{f(x) + g(z)\} \\ & \text{such that } z = Ax - b \end{aligned} \quad (33)$$

and to carry out the minimisations with resp. to x and z in an alternating manner. If there exist closed form representations of the proximal mappings

$$\arg \min_x \left\{ f(x) + \frac{\mu}{2} \|Ax - p_1\|_2^2 \right\} \quad (34)$$

$$\arg \min_z \left\{ g(z) + \frac{\mu}{2} \|z - p_2\|_2^2 \right\} \quad (35)$$

for arbitrary vectors p_1, p_2 , and arbitrary positive μ , then the SB algorithm can be carried out very fast and avoid possible pitfalls that would occur in the formulation of (17) due to the potential non-differentiability. The complete formulation of the SB algorithm to solve (33) is given in Algorithm 1. We remark that the initialisation stated

Algorithm 1: Split Bregman Algorithm for solving (33)

```

1 Initialise  $x^{[0]}$  such that  $0 \in \partial(g)(x^{[0]})$ ,  $z^{[0]} = 0$ ,  $\hat{b} = b$  and  $\mu > 0$  arbitrarily
2 repeat
3   repeat
4      $x^{[k+1]} = \arg \min_x \left\{ f(x) + \frac{\mu}{2} \|Ax - z^{[k]} - \hat{b}\|_2^2 \right\}$ 
5      $z^{[k+1]} = \arg \min_z \left\{ g(z) + \frac{\mu}{2} \|Ax^{[k+1]} - z - \hat{b}\|_2^2 \right\}$ 
6   until convergence
7    $\hat{b} = \hat{b} - (Ax^{[k+1]} - b - z^{[k+1]})$ 
8 until convergence

```

in Algorithm 1 asserts convergence for any choice $\mu > 0$, although the convergence speed may be affected. In practical setups it may be difficult to find an initial value for x such that $0 \in \partial g(x)$. An asset of the SB scheme is that it is still likely to converge with arbitrary initialisations. The two minimising steps in Lines 4 and 5 represent a minimising strategy to solve

$$\arg \min_{x, z} \left\{ f(x) + g(z) + \frac{\mu}{2} \|Ax - z - \hat{b}\|_2^2 \right\}$$

The original authors of the SB algorithm suggest a single sweep in [14]. Our experience suggests that a few more alternating minimisations may be worthwhile for certain applications.

In the forthcoming paragraphs we will show how to fit the formulations from Algorithm 1 to our models. The main difficulty is the fact that our optimisation is being carried out with respect to complex valued (diagonal) matrices.

3.3 Optimising Without Structural Constraints

In this section we derive an iterative strategy to solve (11) and (12). Both models only differ in a weighting matrix and can be handled in the same algorithmic way. We suggest to apply the following mapping

$$f(X) := \frac{1}{2} \|AXA^\top - C\|_F^2 \quad (36)$$

$$g(X) := \lambda \|W \circ X\|_1 \quad (37)$$

where, depending on whether we consider (11) or (12) (resp. (13) in the forthcoming section), either $\lambda = 1$, or W is the matrix having 1 as entry in each position. In addition, we introduce the slack variable by requiring that $X = D$. A straightforward application of the SB algorithm (as presented in (33)) leads now to the following iterative formulation

$$(X, D)^{[k+1]} = \arg \min_{X, D \in \mathbb{R}^{m,m}} \left\{ f(X) + g(D) + \frac{\mu}{2} \|D - X - B^{[k]}\|_F^2 \right\} \quad (38)$$

$$B^{[k+1]} = B^{[k]} - D^{[k+1]} + X^{[k+1]}$$

The minimisation with respect to (X, D) is done in an alternating manner. The optimisation with respect to X reduces to solving the following least squares problem

$$X^{[k+1]} = \arg \min_{X \in \mathbb{R}^{m,m}} \left\{ \frac{1}{2} \|AXA^\top - C\|_F^2 + \frac{\mu}{2} \|X - (D^{[k]} - B^{[k]})\|_F^2 \right\} \quad (39)$$

while the minimisation of D decouples into m^2 optimisations. Indeed, we have to solve

$$D^{[k+1]} = \arg \min_{D \in \mathbb{R}^{m,m}} \left\{ \lambda \|W \circ D\|_1 + \frac{\mu}{2} \|D - X - B^{[k]}\|_F^2 \right\} \quad (40)$$

which can be reduced to component-wise soft shrinkage operations as shown in (31). The necessary optimality conditions for (39) are given by the linear system

$$A^\top (AXA^\top - C) A + \mu (X - D^{[k]} - B^{[k]}) = 0 \quad (41)$$

The minimisation in (40) can be expressed in terms of the soft shrinkage operator and a closed form expression for D_{ij} is given by

$$D_{ij}^{[k+1]} = \text{shrink}_{\frac{\lambda w_{ij}}{\mu}} \left(\text{Re} \left(X_{ij}^{[k+1]} + B_{ij}^{[k]} \right) \right) + \quad (42)$$

$$\iota \cdot \text{shrink}_{\frac{\lambda w_{ij}}{\mu}} \left(\text{Im} \left(X_{ij}^{[k+1]} + B_{ij}^{[k]} \right) \right)$$

Algorithm 2: Split Bregman for solving (11) or (12)

Input: A, C, λ, W, α
Output: Solution of (11) resp. (12).

- 1 Initialise $X = 0, D = 0$ and $B = 0$
- 2 **repeat**
- 3 Set $\hat{X} = X$ and $\hat{D} = D$
- 4 **repeat**
- 5 Set $\bar{X} = \hat{X}$
- 6 **repeat**
- 7 Compute α according to Proposition 5.
- 8 $\bar{X} = \bar{X} - \alpha (A^\top (A\bar{X}A^\top - C)A + \lambda (\bar{X} - \hat{D} + B))$
- 9 **until** convergence towards \bar{X}^*
- 10 $\hat{D} = \text{shrink}_{\frac{\mu W}{\lambda}}(\bar{X}^* + B)$
- 11 **until** convergence towards \hat{X}^* and \hat{D}^*
- 12 Set $X = \hat{X}^*, D = \hat{D}^*$ and $B = B - D + X$
- 13 **until** convergence of X, D and B

We emphasise the importance that W is a real-valued matrix with non-negative entries only. Otherwise it would not be possible to extract its entries and to combine them with the parameter μ in the shrinkage formula.

While the minimisation with respect to D can be carried out very efficiently, this is not necessarily the case for the minimisation with respect to X . Indeed, the occurring linear system has a very large and dense matrix. Solving this system with standard methods from the literature would lead to a strategies that are prohibitive both in terms of memory and run time. Instead, we opt to solve (39) through a simple gradient descent scheme. A single descent step can easily be derived and is given by

$$X^{[k,j+1]} = X^{[k,j]} - \alpha \left(A^\top (AX^{[k,j]}A^\top - C)A + \mu (X^{[k,j]} - D^{[k]} + B^{[k]}) \right) \quad (43)$$

where α is a step size that must be chosen sufficiently small to assert a decrease in energy. We note that the best value for α may be obtained numerically through a simple 1D optimisation and that the best value can also be estimated analytically. We refer to [30] for an overview of several efficient numerical strategies and to Proposition 5 for the computation of the optimal step size. The complete algorithm is given in Algorithm 2. Let us also remark that the two nested inner loops in the algorithm do not need to be carried out until full convergence. For many applications already a single iteration step is usually enough: In [14] the authors suggested to use a single Gauß-Seidel iteration to solve the occurring linear system. The influence of inaccuracies in the Bregman iteration has also been studied in [49]. In general the SB scheme is very robust against inaccuracies. The benefit of reducing the number of iterations will usually yield significant increase in performance. The first two nested loops in Algorithm 2 correspond in a one-to-one manner to the loops in Algorithm 1. The third loop beginning in Line 6 represents the gradient descent scheme explained in (43). It corresponds to the minimisation with respect to x in Line 4 in Algorithm 1.

The following proposition states the optimal value for the gradient descent scheme, in the sense that in each iteration the decrease in the energy is maximal.

Proposition 5 *The optimal value for the step size α for the gradient descent scheme from (43) is given by:*

$$\alpha = \frac{\|G(X^{[k,j]})\|_F^2}{\|AG(X^{[k,j]})A^\top\|_F^2 + \mu \|G(X^{[k,j]})\|_F^2} \quad (44)$$

where G is the first order derivative with respect to X of the cost functional

$$E(X) := \frac{1}{2} \|AXA^\top - C\|_F^2 + \frac{\mu}{2} \|D^{[k]} - X - B^{[k]}\|_F^2 \quad (45)$$

Proof Let G be the gradient of our energy E , i.e.

$$G(X) = A^\top (A X A^\top - C)A + \mu (X - D^{[k]} + B^{[k]}) \quad (46)$$

In order to maximise the energy decrease in a single gradient descent step we wish to solve

$$\arg \min_{\alpha} \left\{ E \left(X^{[k,j]} - \alpha G(X^{[k,j]}) \right) \right\} \quad (47)$$

Using the definition of E , we are led to the following expression, that needs to be minimised with respect to α :

$$\frac{1}{2} \left\| \alpha AG(X^{[k,j]})A^\top - (AXA^\top - C) \right\|_F^2 + \frac{\mu}{2} \left\| \alpha G(X^{[k,j]}) - (X^{[k,j]} - D + B^{(k)}) \right\|_F^2 \quad (48)$$

Using the properties of the matrix scalar product, it follows that the latter expression can be rewritten as

$$\begin{aligned} E(X^{[k,j]}) + \alpha^2 \left(\frac{1}{2} \|AG(X^{[k,j]})A^\top\|_F^2 + \frac{\mu}{2} \|G(X^{[k,j]})\|_F^2 \right) \\ - \alpha \langle AG(X^{[k,j]})A^\top, AX^{[k,j]}A^\top - C \rangle - \alpha \mu \langle G(X^{[k,j]}), X^{[k,j]} - D + B^{(k)} \rangle \end{aligned} \quad (49)$$

Thus, the optimal α is given by

$$\begin{aligned} \alpha &= \frac{\langle AG(X^{[k,j]})A^\top, AX^{[k,j]}A^\top - C \rangle + \mu \langle G(X^{[k,j]}), X^{[k,j]} - D + B^{(k)} \rangle}{\|AG(X^{[k,j]})A^\top\|_F^2 + \mu \|G(X^{[k,j]})\|_F^2} \\ &= \frac{\langle G(X^{[k,j]}), A^\top AX^{[k,j]}A^\top A - A^\top CA + \mu (X^{[k,j]} - D + B^{(k)}) \rangle}{\|AG(X^{[k,j]})A^\top\|_F^2 + \mu \|G(X^{[k,j]})\|_F^2} \\ &= \frac{\|G(X^{[k,j]})\|_F^2}{\|AG(X^{[k,j]})A^\top\|_F^2 + \mu \|G(X^{[k,j]})\|_F^2} \end{aligned} \quad (50)$$

□

3.4 Structured Split Bregman Model

We now consider our third model from (13). Our approach is similar to the approach presented in Section 3.3. This time we additionally enforce a diagonal structure of X by restricting the optimisation with respect to X to matrices having the desired structure. Our functional takes the form

$$\arg \min_X \left\{ \frac{1}{2} \|AXA^\top - C\|_F^2 + \mu \|X\|_1 \right\} \quad (51)$$

where μ is again a sparsity inducing regularisation weight. We abstain from using a more fine-grained weighting matrix W since we do not have any prior information on the sparsity pattern. Nevertheless, we remark that such a modification would be possible.

This time the minimisation is not done over the set of all matrices but merely over the set of diagonal matrices. Our consideration is that the reduced search space improves performance and yields optima closer to our expectations.

The minimisation is done with the help of the SB approach. Similarly as before, we introduce a dummy variable D and obtain

$$(X, D)^{[k+1]} = \arg \min_{X, D} \left\{ \frac{1}{2} \|AXA^\top - C\|_F^2 + \mu \|D\|_1 + \frac{\lambda}{2} \|D - X - B^{[k]}\|_F^2 \right\} \quad (52)$$

$$B^{[k+1]} = B^{[k]} - D^{[k+1]} + X^{[k+1]}$$

The minimisation with respect to X and D is done in an alternating manner. The optimisation with respect to D is identical to the previous section where we set $W_{ij} = 1$ for all i and j . If D is initialised as a diagonal matrix, then no special consideration need to be done to preserve this structure provided that X is initialised as a diagonal matrix, too. The minimisation with respect to X cannot be reduced to a linear system anymore since we have introduced a structure preserving constraint. Nevertheless, an explicit form of the gradient of the energy to be minimised is available. Using the findings from Proposition 4 we conclude that the gradient of our cost function in (52) with respect to the *diagonal matrix* X can be expressed as

$$\left(A^\top (AXA^\top - C) A \right) \circ I + \lambda \left(X - D^{(k)} + B^{(k)} \right) \circ I \quad (53)$$

The two Hadamard products with the identity matrix in the previous formula stem from the fact that we compute the derivative with respect to diagonal matrices. They ensure that the diagonal structure of the matrix X is preserved throughout the whole iterative process.

The full algorithm is given in Algorithm 3. Also here it is possible to optimise the step size for the gradient descent scheme in a similar manner as in Proposition 5.

3.5 Post Processing Steps

Numerical experiments show that sometimes it is necessary to post-process the obtained signals to remedy certain deficiencies that stem from the experimental nature

Algorithm 3: Split Bregman for solving (13)

Input: A, C, α, μ
Output: Solution of (13).

- 1 Initialise $X = 0, D = 0$ and $B = 0$
- 2 **repeat**
- 3 Set $\hat{X} = X$ and $\hat{D} = D$
- 4 **repeat**
- 5 Set $\bar{X} = \hat{X}$
- 6 **repeat**
- 7 Compute optimal α .
- 8 $\bar{X} = \bar{X} - \alpha (A^\top (A\bar{X}A^\top - C)A + \lambda (\bar{X} - \hat{D} + B)) \circ I$
- 9 **until** convergence towards \bar{X}^*
- 10 $\hat{D} = \text{shrink}_{\frac{\mu}{\lambda}}(\hat{X} + B)$
- 11 **until** convergence towards \hat{X}^* and \hat{D}^*
- 12 Set $X = \hat{X}^*, D = \hat{D}^*$ and $B = B - D + X$
- 13 **until** convergence of X, D and B

of the physical setup. The following simple strategy has proven to be efficient and reliable. We formulate it for the results obtained from Section 3.4, but an analogous application to the results obtained with the algorithm from Section 3.3 are clearly possible, too.

1. Remap the diagonal entries from the matrix X to their actual positions in 2D (resp. 3D) space. Thus, we obtain tuples (i, j, x_{ij}) .
2. Apply a k-means clustering [26] approach to partition the tuples into distinct sets. If the exact number of clusters is unknown then it can either be estimated with the strategy from [38] or by using other popular strategies commonly used in the clustering context, such as silhouette coefficients [36] or GAP statistics [43].
3. Use the centroid position of each cluster as source position.
4. Sum up all source strengths from a cluster to obtain the corresponding source strength.

Our experiments have shown that the sum of all entries of our computed solutions always coincided with the sum of all entries of the ground truth solution. This observation motivates our suggestion to sum all entries of a cluster. Further knowledge on the solution could also be introduced into the post processing. As such, it would for example be possible to introduce a minimal distance between two peaks. This could help in handling noise, as it tends to cluster small fluctuations around a single large peak.




Fig. 2 Graphical representation of the ground truth of our simulated experiments. *Left:* The entries from the diagonal matrix X with their corresponding indices and values. Only the indicated entries are non-zero. *Right:* The corresponding spatial distribution.

4 Numerical Evaluation

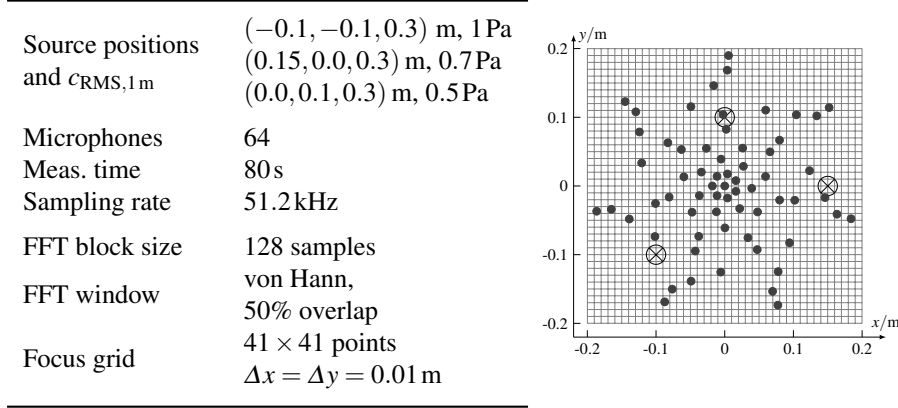
We evaluate our models and show the benefits of the additional post processing. For each tested model we also discuss the influence of noisy data and inaccuracies in the measurements. The findings are depicted in their corresponding sections.

Our tests use different data setups. The considered data stems from simulated experiments and therefore, an exact ground truth is to our avail. A detailed presentation of the considered setup is given in Section 4.1. In our setup we have $A \in \mathbb{C}^{64,1681}$ and $C \in \mathbb{C}^{64,64}$, such that X must be a diagonal matrix in $\mathbb{C}^{1681,1681}$. The exact solution is supposed to contain 3 sources with different strengths. The corresponding diagonal matrix X_{sol} has a non-zero entry in the positions 421, 851 and 1456 of its main diagonal. The corresponding values are 0.1422, 0.0392 and 0.0682 respectively. A plot visualising the solution in form of the diagonal entries from the matrix X as well as their spatial distribution is given in Fig. 2.

Our numerical tests cover the commonly occurring use-cases where different frequencies and noise corruption must be handled. Data corruption can occur as noise applied onto the recorded signals or as inaccuracies in the measurement of the position. The latter case is much more difficult to handle since there is no clear way to deduce the correct location from faulty data. This problem is severely ill-posed. We remark, that the physical setup is encoded in the matrix C . The matrix A is given by the sound propagation model detailed in Section 1.

We will refrain from using any error measure, such as the popular ℓ_2 - or ℓ_1 -distance since they do not allow to reflect the sparsity requirements in a convincing manner. Neither of these distances reflect the concrete physical properties of the underlying problem and further, they cannot distinguish between an additional source with a certain strength and a wrongly estimated source with the same deviation. In both cases the distance to the ground truth would be the same. Therefore, we will only provide graphical representations of the solution. The plots accurately visualise the determined number of sources. Furthermore, since our setups are very sparse, it is easy to compare the strengths of the individual sources to the ground truth.

Fig. 3 *Left:* Data generation parameters. *Right:* Distribution of the 64 microphones (filled circles) with the simulated sources (crossed circles). Also shown is the discretised grid of focus points. The array is positioned 0.3 m above the sources, parallel to the focus grid.



In the following we evaluate our model from (12) and from (13) and compare our results to reference solutions obtained with well established methods from the literature. We do not evaluate the model from (11) explicitly, since it is a special case of the formulation in (12) where all entries in the weighting matrix take the same value. Furthermore, (11) provides no mechanism to favour a diagonal structure of the solution X . Therefore, we expect it to perform significantly worse than the other two strategies in the experiments considered in this work. Nevertheless, we point out that (11) is probably a valuable and competitive strategy when the sources are not assumed to be uncorrelated. In this case non-zero entries outside of the main diagonal are to be expected.

4.1 Basic Setup

The exemplary cross-spectral matrix C of microphone signals used for this study is synthesised from simulated data, following a typical data processing used with microphone array measurements for evaluation in the frequency domain.

First, microphone time signals are simulated by calculating the sound pressures caused by three monopole sources distributed at different positions in one plane parallel to the array. These sources emit uncorrelated Gaussian white noise with differing source strengths, defined by the Root Mean Square (RMS) value of the sound pressure in 1 m distance from the respective source. Data generation parameters and the array geometry with the relative source positions are specified in Fig. 3. Following Welch's method [44], the time signals are cut into K overlapping blocks of equal length, onto which a Fast Fourier Transform (FFT) is applied. For each block, the cross-spectra between the microphone channels are calculated from the vector of complex sound pressures c_k . Finally, the cross-spectral matrix is estimated

by averaging the cross-spectra:

$$C = \frac{1}{K} \sum_k c_k c_k^\top, \quad k = 1, \dots, K \quad (54)$$

The assumed possible source locations form a square grid of regularly spaced focus points, contain the actual source positions, and are arranged in a plane parallel to the array. The extension of the focus grid is of the same order of magnitude as that of the array (see Fig. 3). With a sampling rate of 51.2 kHz and a FFT block size of 128 we can compute 63 cross-spectral matrices from the microphone signals. Each matrix represents a discrete frequency band with band centre frequencies between 400 Hz and 25.2 kHz and with a band width of 400 Hz. All experiments in this paper use a cross-spectral matrix with a band centre frequency of 19.2 kHz.

The noisy data is generated by corrupting the microphone signals directly. Each microphone is corrupted individually. The noise applied to the signals is uncorrelated and has a relative strength of 3.33, 4.76 and 6.25 respectively, when compared to the signal strength. In particular, the noise is stronger than the actual signal.

Additionally, a data set with erroneous microphone positions was generated. For this, the nominal position of each microphone was disturbed randomly within the array plane. The average deviation from the true position is 0.009 m. The array has an aperture of 0.4 m.

4.2 Using the Weighted Model from Equation (12)

In this section we analyse the performance of our model stated in (12). Our first test uses optimal uncorrupted data. We set the number of Bregman iterations, alternating minimisations and gradient descent steps to 75/3/20. The weighting matrix has the value 1 on its main diagonal and 10^6 in all other positions. The Bregman regularisation weight has the value 10^5 . With these parameters the run time of a pure Matlab implementation was about 20 minutes for a single frequency band. Our algorithms are always initialised with a solution having the value 1 at all positions. The run time of the post processing is negligible. The k-means algorithm always converged within a few iterations and had run times below one second.

4.2.1 Using Perfect Data

The results obtained from the perfect data are visualised in Fig. 4. As we can see, the reconstruction is almost perfect. The diagonal has very few non-zero entries, that cluster around the exact solutions. All the other entries are smaller than 10^{-5} in magnitude and have been excluded from the plot to ease the visualisation. After the post-processing step with our k-means approach we obtain the reconstruction in the right plot in Fig. 4.




Fig. 4 Results obtained from Algorithm 2 for solving (12) for the noise-free setup. *Left*: Diagonal entries returned by the algorithm. The grey discs represent the solution without post processing, whereas the grey squares denote the post-processed solution. The dashed vertical lines indicate the cluster boundaries obtained by our k-means algorithm. We see three distinct clusters of entries centred around the exact solution (marked in black). *Right*: The corresponding graphical representation on the 2D Grid. The filled grey shapes indicate the positions of the non-zero entries from the main diagonal. Each shape represents one of the clusters returned by our post-processing strategy. The crosses and the grey labels mark the position and strength of the computed centres of the clusters. The white labels and circles depict the ground truth solution.

4.2.2 Using Noisy Data

Similar findings to those from the previous paragraph are also obtainable when the data is corrupted by noise. The results are given in Fig. 5. They show that our approach is robust to noise. The obtained solutions, especially after the post-processing, hardly differ from those from the perfect setup.

4.2.3 Using Data with Positional Errors

Finally, the data with positional errors yields the results presented in Fig. 6. Without the k-means post processing the solution is hardly usable. However, with the additional clustering step we obtain a rather good reconstruction. Let us emphasise that the handling of the data sets with errors in the positional data is a severely ill posed problem for which no satisfactory solution exists in the literature so far.

4.3 Using the Model with Structural Constraints from Equation (13)

In this section we want to demonstrate the benefits of integrating structural constraints into the model and optimisation process. We test our Algorithm 3 on the same data as in the previous section. Applying Algorithm 3 with 75 outer iterations, 4 alternating minimisation steps and 10 gradient descent steps.




Fig. 5 Results obtained from Algorithm 2 for solving (12) for the noisy setup. *Left*: Diagonal entries returned by the algorithm. The grey discs represent the solution without post processing, whereas the grey squares denote the post-processed solution. The dashed vertical lines indicate the cluster boundaries obtained by our k-means algorithm. We see three distinct clusters of entries centred around the exact solution (marked in black). *Right*: The corresponding graphical representation on the 2D Grid. The filled grey shapes indicate the positions of the non-zero entries from the main diagonal. Each shape represents one of the clusters returned by our post-processing strategy. The crosses and the grey labels mark the position and strength of the computed centres of the clusters. The white labels and circles depict the ground truth solution. As we can see, our algorithm behaves well, even in the presence of noise. The results are almost identical to those from Fig. 4




Fig. 6 Graphical representation of the diagonal matrix using data with erroneous positional information as returned by Algorithm 2. *Left*: A straightforward application yields an undesirable outcome. There are many non-zero entries along the main diagonal. The grey discs represent the solution without post processing, whereas the grey squares denote the post-processed solution. The dashed vertical lines indicate the cluster boundaries obtained by our k-means algorithm. After applying the post processing steps we obtain a convincing localisation. The black circles denote the exact ground truth. *Right*: The corresponding graphical representation on the 2D Grid. The filled grey shapes indicate the positions of the non-zero entries from the main diagonal. Each shape represents one of the clusters returned by our post-processing strategy. The crosses and the grey labels mark the position and strength of the computed centres of the clusters. The white labels and circles depict the ground truth solution.




Fig. 7 Results obtained from Algorithm 3 for solving (13) for the noise-free setup. *Left*: Diagonal entries returned by the algorithm. The grey discs represent the solution without post processing, whereas the grey squares denote the post-processed solution. The dashed vertical lines indicate the cluster boundaries obtained by our k-means algorithm. Even without the post processing the solution is nearly indistinguishable from the exact solution (marked in black). *Right*: The corresponding graphical representation on the 2D Grid. The filled grey shapes indicate the positions of the non-zero entries from the main diagonal. Each shape represents one of the clusters returned by our post-processing strategy. The crosses and the grey labels mark the position and strength of the computed centres of the clusters. The white labels and circles depict the ground truth solution.

4.3.1 Using Perfect Data

The results obtained from the perfect data are visualised in Fig. 7. Here, the sparsity weight was set to 10 and the Bregman regularisation parameter was set to 10^4 . As we can see, the obtained result is almost identical to the reference ground truth.

4.3.2 Using Noisy Data

Figure 8 depicts our reconstruction from the noise corrupted signal. As we can see, our approach is robust and still yields almost ideal results. The parameters were exactly the same as in the noise free case. Our experiments suggest that the method works reliably, for almost any reasonable parameter setting. We consider this a further asset of our strategy.

4.3.3 Using Data with Positional Errors

The data set with the encoded erroneous positional information has proven to be the most difficult to handle. Executing our method with 75 Bregman iterations, 6 alternating minimisations and 50 gradient descent steps yields the results shown in Fig. 9. Other parameter choices did not yield significantly better results. As we can see, the entries on the main diagonal of the matrix are not as sparse as before and they are off by almost a factor 10. In order to remedy the situation, we apply our k-means based post-processing. The resulting clusters, and their centroid positions are presented in Fig. 9. The imaginary part of every entry is smaller than 10^{-5} in magnitude. Also, all entries smaller in absolute value than 10^{-3} have been omitted




Fig. 8 Results obtained from Algorithm 3 for solving (13) for the noisy data setup. Again, our approach shows a very favourable behaviour in presence of noise. Our results are identical to those from the noise free test case.




Fig. 9 Graphical representation of the of the diagonal matrix using data with erroneous positional information as returned by Algorithm 3. *Left:* A straightforward application yields an undesirable outcome. There are many non-zero entries along the main diagonal. Also here, the post processing improves the results significantly. *Right:* The corresponding graphical representation on the 2D Grid. The filled grey shapes indicate the positions of the non-zero entries from the main diagonal. Each shape represents one of the clusters returned by our post-processing strategy. The crosses and the grey labels mark the position and strength of the computed centres of the clusters. The white labels and circles depict the ground truth solution.

to improve the visualisation. Clearly, our result consists of 3 distinct clusters which are well localised. The position of each cluster coincides with the coordinates of one peak from the ground truth. However, the scaling does not match. Yet, it is interesting to note that the sum of all entries in a single cluster almost gives the sought solution value.

4.4 Comparison to Other Methods from the Literature

Let us shortly compare the output of our schemes to well established methods from the literature. To this end we have evaluated our experimental data with the Clean




Fig. 10 Solution from perfect data with the CMF method. Except for two very weak additional sources detected in the vicinity of the strongest sources the method finds a near perfect solution. The positions of the three strongest sources coincide with the ground truth, only the intensities are slightly off. *Left*: Graphical representation of the main diagonal. *Right*: Corresponding spatial distribution. Crosses mark the obtained solution whereas the labelled circles denote the ground truth. The labels of the obtained solution have been omitted to avoid unnecessary cluttering of the visualisation.

based on Source Coherence (Clean-SC) [41] and Covariance Matrix Fitting (CMF) [2, 48] strategies. Clean-SC is an extension of the venerable CLEAN algorithm for deconvolution [22]. The CMF method is similar in spirit as our optimisation models from (11). We use a soft constraint on the sparsity, whereas CMF uses a hard constraint.

The methods have been implemented and used “as-is” without any additional post or preprocessing of the data. We remark that the considered methods are tailored towards yielding sparse solutions. As a consequence, post processing steps, such as our clustering approach, make little sense in this context.

4.4.1 Evaluation of the Covariance Matrix Fitting Method

Using the CMF method we obtain the results presented in Figs. 10, 11 and 12. The figures represent the use-cases with perfect and noisy data as well as with erroneously encoded positions respectively. As we can observe in Fig. 10, the CMF method fails to accurately detect the correct number of sources. It adds two additional sources (with almost negligible strength). Nevertheless, the three largest peaks coincide almost perfectly with the ground truth solution from Fig. 2. In the case of noisy data in Fig. 11, the CMF approach works equally well. Position and magnitude of the three largest peaks are very accurate. Finally, the method completely fails if the positions are erroneously encoded (c.f. Fig. 12). Not only the position, but also the number and the strength of the signals are incorrectly detected.

4.4.2 Evaluation of the Clean-SC Method

Using the Clean-SC method we obtain the results presented in Figs. 13, 14 and 15. The figures represent the use-cases with perfect and noisy data, as well as with erroneously encoded positions respectively. In presence of perfect data, the Clean-SC method




Fig. 11 Solution from noisy data with the CMF method. Except for two very weak additional sources detected in the vicinity of the strongest sources the method finds a near perfect solution. The solution of the CMF algorithm is almost unaffected by the addition of noise. *Left:* Graphical representation of the main diagonal. *Right:* Corresponding spatial distribution. Crosses mark the obtained solution whereas the labelled circles denote the ground truth. The labels of the obtained solution have been omitted to avoid unnecessary cluttering of the visualisation.




Fig. 12 Solution from data with erroneous positions with the CMF Method. In this case the algorithm breaks completely down. Not only the positions and number of sources are wrong, but also the signal strengths are in a completely wrong range. *Left:* Graphical representation of the main diagonal. *Right:* Corresponding spatial distribution. Crosses mark the obtained solution whereas the labelled circles denote the ground truth. The labels of the obtained solution have been omitted to avoid unnecessary cluttering of the visualisation.

works extraordinarily well and even outperforms the CMF algorithm. The yielded results are nearly perfect, see Fig. 13. On the other hand, if the data is corrupted by noise, then the Clean-SC strategy cannot compete with CMF. The method yields too many sources as can be seen in Fig. 14. Furthermore, the strength of certain sources is completely wrong. The largest peak in Fig. 14 is nearly three times larger than it should. In presence of positional errors the method performs similarly to CMF, the findings are little convincing, as can be observed in Fig. 15.




Fig. 13 Solution from perfect data with the Clean-SC Method. The number of source, their position and their strength are detected with a very high accuracy. *Left:* Graphical representation of the main diagonal. *Right:* Corresponding spatial distribution. Crosses mark the obtained solution whereas the labelled circles denote the ground truth. The labels of the obtained solution have been omitted to avoid unnecessary cluttering of the visualisation.




Fig. 14 Solution from noisy data with the Clean-SC Method. The algorithm suffers more from noisy data than CMF. The number of sources is wrongly estimated and the magnitude of the two largest peaks is completely wrong. The ticks on the y axis denote the correct magnitude. As we can see, the third peak is about 3 times larger than it should. *Left:* Graphical representation of the main diagonal. *Right:* Corresponding spatial distribution. Crosses mark the obtained solution whereas the labelled circles denote the ground truth. The labels of the obtained solution have been omitted to avoid unnecessary cluttering of the visualisation.

4.5 Conclusion

The following conclusions can be drawn from the evaluation.

1. State-of-the-art methods can handle perfect data as well as noisy data quite well.
2. Our algorithms yield competitive results for perfect and noisy data. One advantage of our method is that the number of sources is always accurate.
3. CMF and Clean-SC cannot handle data sets with wrongly encoded positions at all.
4. Our algorithms do not yield perfect solutions for wrongly encoded positions either, nevertheless the quality is by far superior to the findings from CMF and Clean-SC.




Fig. 15 Solution from data with erroneous positions with the Clean-SC Method. Clean-SC behaves similarly as CMF. *Left:* Graphical representation of the main diagonal. *Right:* Corresponding spatial distribution. Crosses mark the obtained solution whereas the labelled circles denote the ground truth. The labels of the obtained solution have been omitted to avoid unnecessary cluttering of the visualisation.

In the total, our new algorithms provide accurate and more robust alternatives compared to actual state-of-the-art methods.

5 Summary and Outlook

We have presented three approaches to retrieve a sparse set of sound source locations from acoustic measurements. Our novel algorithms apply findings from matrix differential calculus to sparsity favouring convex optimisation models, which we have applied onto complex matrix valued settings. Our efforts to integrate the structural constraints, such as the shape of the matrix to be optimised into the optimisation problem, have been rewarded by very convincing results. Not only did we benefit from the reduced memory footprint and faster run times, but also the obtained results are much more accurate than in the unconstrained framework. Empirical studies show that our algorithms are competitive to state-of-the-art methods from the literature for perfect and noisy data. In presence of ill-posed setups, we benefit from our post processing strategies to obtain reasonable results.

In the future, we would like to provide further improvements to the accuracy of the results. Especially the handling of inaccuracies in the recording of the positions will be of interest to us. Further investigations on acceleration methods for Bregman schemes will also be a topic of our research. We are confident that a significant reduction of the run time is possible.

References

1. Bauschke, H.H., Borwein, J.M.: Legendre functions and the method of random Bregman projections. *Journal of Convex Analysis* **4**(1), 27–67 (1997)
2. Blacodon, D., Elias, G.: Level estimation of extended acoustic sources using a parametric method. *Journal of Aircraft* **41**(6), 1360–1369 (2004)

3. Bregman, L.M.: The relaxation method of finding the common point of convex sets and its application to the solution of problems in convex programming. *USSR Computational Mathematics and Mathematical Physics* **7**(3), 200–217 (1967)
4. Brune, C., Sawatzky, A., Burger, M.: Primal and dual Bregman methods with application to optical nanoscopy. *International Journal of Computer Vision* **92**, 211–229 (2011)
5. Burger, M., Gilboa, G., Osher, S., Xu, J.: Nonlinear inverse scale space methods. *Communications in Mathematical Sciences* **4**(1), 175–208 (2006)
6. Burger, M., Osher, S., Xu, J., Gilboa, G.: Nonlinear inverse scale space methods for image restoration. In: N. Paragios, O.D. Faugeras, T. Chan, C. Schnörr (eds.) *Variational, Geometric, and Level Set Methods in Computer Vision (VLSM)*. Third international workshop, *Lecture Notes in Computer Science*, vol. 3752, pp. 25–36. Springer, Beijing, China (2005)
7. Burger, M., Resmerita, E., He, L.: Error estimation for Bregman iterations and inverse scale space methods. *Computing* **81**(2–3), 109–135 (2007)
8. Cai, J.F., Osher, S., Shen, Z.: Linearized Bregman iterations for compressed sensing. *Mathematics of Computation* **78**(267), 1515–1536 (2009)
9. Cai, J.F., Osher, S., Shen, Z.: Split Bregman methods and frame based image restoration. *Multiscale Modeling & Simulation* **8**(2), 337–369 (2009)
10. Chambolle, A., Pock, T.: A first-order primal-dual algorithm for convex problems with applications to imaging. *Journal of Mathematical Imaging and Vision* **40**, 120–145 (2011)
11. Geiger, C., Kanzow, C.: *Theorie und Numerik restringierter Optimierungsaufgaben*. Springer-Lehrbuch Masterclass. Springer (2002)
12. Ghadimi, E., Teixeira, A., Shames, I., Johansson, M.: Optimal parameter selection for the alternating direction method of multipliers (admm): Quadratic problems. *IEEE Trans. Automat. Contr.* **60**(3), 644–658 (2015)
13. Goldstein, T., O'Donoghue, B., Setzer, S., Baraniuk, R.: Fast alternating direction optimization methods. *SIAM J. Imaging Sci.* **7**(3), 1588–1623 (2014)
14. Goldstein, T., Osher, S.: The split gBregman method for ℓ_1 regularized problems. *SIAM Journal on Imaging Sciences* **2**(2), 323–343 (2009)
15. Goldstein, T., Taylor, G., Barabin, K., Sayre, K.: Unwrapping admm: Efficient distributed computing via transpose reduction. Arxiv Preprint 1504.02147, University of Maryland, College Park and United States Naval Academy, Annapolis (2015)
16. He, L., Chang, T.C., Osher, S., Fang, T., Speier, P.: MR image reconstruction by using the iterative refinement method and nonlinear inverse scale space methods. UCLA CAM Report 06-35, University of California, Los Angeles (2006)
17. Herold, G., Sarradj, E.: Preliminary benchmarking of microphone array methods. In: *Proceedings, Berlin Beamforming Conference, 2014*. Gesellschaft zur Förderung angewandter Informatik (2014)
18. Herold, G., Sarradj, E.: An approach to estimate the reliability of microphone array methods. In: *21st AIAA/CEAS Aeroacoustics Conference* (2015)
19. Herold, G., Sarradj, E., Geyer, T.: Covariance matrix fitting for aeroacoustic application. In: *Fortschritte der Akustik - AIA-DAGA 2013*, pp. 1926–1928 (2013)
20. Horn, R.A., Johnson, C.R.: *Matrix Analysis*. Cambridge University Press (1990)
21. Horn, R.A., Johnson, C.R.: *Topics in Matrix Analysis*. Cambridge University Press (1994)
22. Högbom, J.A.: Aperture synthesis with a non-regular distribution of interferometer baselines. *Astronomy and Astrophysics Supplement* **15**, 417–426 (1974)
23. Kadkhodaie, M., Christakopoulou, K., Sanjabi, M., Banerjee, A.: Accelerated alternating direction method of multipliers. In: *KDD '15 Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pp. 497–506. ACM New York (2015)
24. Lin, C.J.: Projected gradient methods for nonnegative matrix factorization. *Neural Computation* **19**(10), 2756–2779 (2007)
25. Magnus, J.R., Neudecker, H.: *Matrix Differential Calculus with Applications in Statistics and Econometrics*. Wiley Series in Probability and Statistics. John Wiley & Sons (2007)
26. McQueen, J.B.: Some methods for classification and analysis of multivariate observations. In: *Proceedings of 5th Berkeley Symposium on Mathematical Statistics and Probability*, vol. 1, pp. 281–297. University of California Press (1967)
27. Million, E.: The Hadamard product (2007). Available online: <http://buzzard.ups.edu/courses/2007spring/projects/million-paper.pdf>
28. Minka, T.P.: Old and new matrix algebra useful for statistics (2000). Available online: <http://research.microsoft.com/en-us/um/people/minka/papers/matrix/>

29. Moreau, J.J.: Proximité et dualité dans un espace hilbertien. *Bulletin de la Société Mathématique de France* **93**, 273–299 (1965)
30. Nocedal, J., Wright, S.J.: *Numerical Optimization*. Springer Series in Operations Research and Financial Engineering. Springer-Verlag New York (2006)
31. Osher, S., Burger, M., Goldfarb, D., Xu, J., Yin, W.: An iterative regularization method for total variation-based image restoration. *Multiscale Modeling and Simulation* **4**(2), 460–489 (2005)
32. Osher, S., Mao, Y., Dong, B., Yin, W.: Fast linearized Bregman iteration for compressive sensing and sparse denoising. *Communications in Mathematical Sciences* **8**(1), 93–111 (2010)
33. Petersen, K.B., Pedersen, M.S.: *The matrix cookbook* (2012). Available online: <https://www.math.uwaterloo.ca/~hwolkowi/matrixcookbook.pdf>
34. Pock, T., Chambolle, A.: Diagonal preconditioning for first order primal dual algorithms in convex optimization. In: 2011 IEEE International Conference on Computer Vision (ICCV), pp. 1762–1769. IEEE (2011)
35. Pollock, D.S.G.: Tensor products and matrix differential calculus. *Linear Algebra and Its Applications* **67**, 169–193 (1985)
36. Rousseeuw, P.J.: Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Computational and Applied Mathematics* **20**, 53–65 (1987)
37. Rudin, L.I., Osher, S., Fatemi, E.: Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena* **60**(1–4), 259–268 (1992)
38. Sarraji, E.: A fast signal subspace approach for the determination of absolute levels from phased microphone array measurements. *Journal of Sound and Vibration* **329**, 1553–1569 (2010)
39. Setzer, S.: Split Bregman algorithm, Douglas-Rachford splitting and frame shrinkage. In: X.C. Tai, K. Mörken, M. Lysaker, K.A. Lie (eds.) *Scale Space and Variational Methods in Computer Vision, Lecture Notes in Computer Science*, vol. 5567, pp. 464–476. Springer (2009)
40. Setzer, S.: Operator splittings, Bregman methods and frame shrinkage in image processing. *International Journal of Computer Vision* **92**, 265–280 (2010)
41. Sijtsma, P.: CLEAN based on spatial source coherence. *International Journal of Aeroacoustics* **6**(4), 357–374 (2007)
42. Tibshirani, R.: Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society. Series B (Methodology)* **58**(1), 267–288 (1996)
43. Tibshirani, R., Walther, G., Hastie, T.: Estimating the number of clusters in a data set via the gap statistic. *J. R. Statistic. Soc. B* **63**(2), 411–423 (2001)
44. Welch, P.D.: The Use of Fast Fourier Transform for the Estimation of Power Spectra: A Method Based on Time Averaging Over Short, Modified Periodograms. *IEEE Transactions on Audio and Electroacoustics* **15**(2), 70–73 (1967)
45. Xu, J., Osher, S.: Iterative regularization and nonlinear inverse scale space applied to wavelet-based denoising. *IEEE Transactions on Image Processing* **16**(2), 534–544 (2006)
46. Xu, Y., Yin, W.: A block coordinate descent method for multi-convex optimization with applications to nonnegative tensor factorization and completion. *Rice CAAM Technical Report TR12-15*, Rice University (2012)
47. Xu, Y., Yin, W., Wen, Z., Zhang, Y.: An alternating direction algorithm for matrix completion with nonnegative factors. *Frontiers of Mathematics in China* **7**(2), 365–384 (2012)
48. Yardibi, T., Li, J., Stoica, P., Cattafesta, L.N.: Sparsity constrained deconvolution approaches for acoustic source mapping. *The Journal of the Acoustic Society of America* **123**(5), 2631–2642 (2008)
49. Yin, W., Osher, S.: Error forgetting of Bregman iteration. *Journal of Scientific Computing* **54**(2), 684–695 (2013)
50. Yin, W., Osher, S., Goldfarb, D., Darbon, J.: Bregman iterative algorithms for ℓ_1 -minimization with applications to compressed sensing. *SIAM Journal on Imaging Sciences* **1**(1), 143–168 (2007)
51. Zhang, X., Burger, M., Bresson, X., Osher, S.: Bregmanized nonlocal regularization for deconvolution and sparse reconstruction. *UCLA CAM Report 09-03*, University of California, Los Angeles (2009)